


Artificial Cognitive Systems

David Vernon
Carnegie Mellon University Africa

vernon@cmu.edu
www.vernon.eu

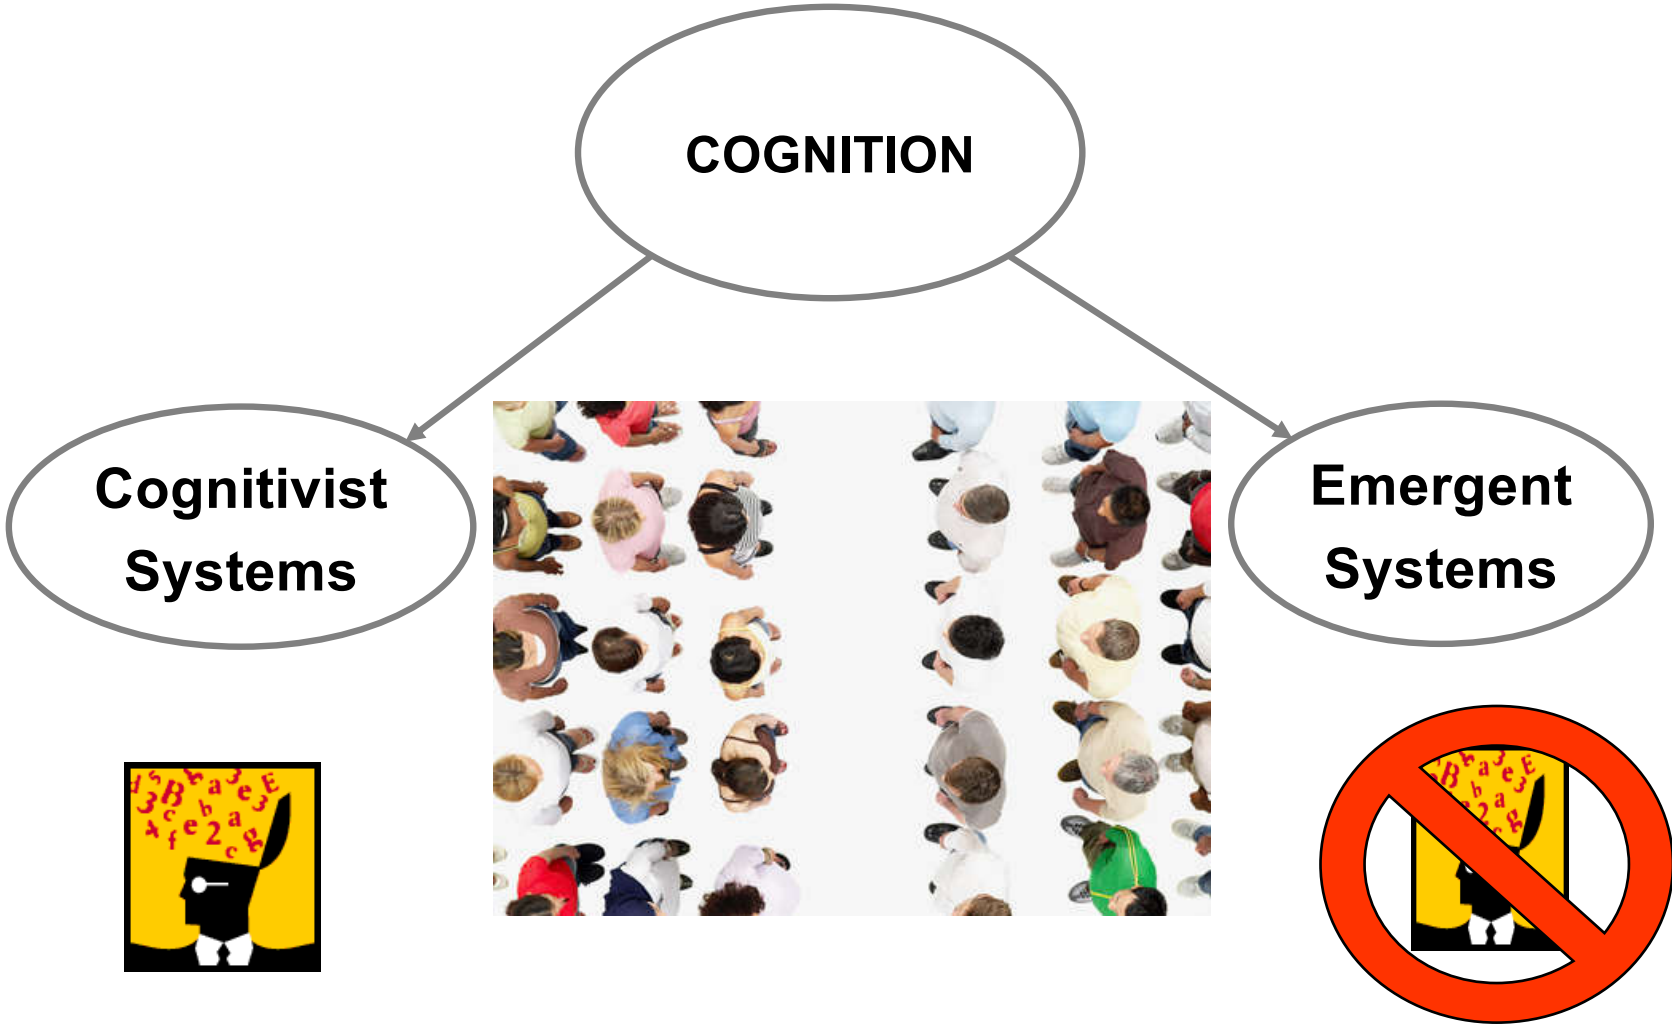
Knowledge & Representation



Topic Overview

- The duality of memory and knowledge
- Representation and anti-representation
- The symbol grounding problem
- Joint perceptuo-motor representations
- Acquiring and sharing knowledge

Knowledge and Representation



The Duality of Memory and Knowledge

Memory and knowledge are intimately related

- Declarative, procedural, episodic, and semantic knowledge
- But there are **hidden assumptions** about the nature of that knowledge

The Duality of Memory and Knowledge

Cognitivist paradigm

- Two tacit assumptions
 1. Knowledge is the content that **complements the cognitive architecture**
 2. Very often, knowledge, even procedural knowledge, is assumed to be **symbolic**

The Duality of Memory and Knowledge

Emergent paradigm

- Memory viewed as both **content** and **process**: as a mechanism for prediction and recollection
- Knowledge is the **manifestation of that process**: it is what emerges when memory works effectively
- Knowledge and memory as complementary aspects of the same thing
- Representation of that knowledge: **symbolic, non-symbolic** (sub-symbolic)

Representation and Anti-Representation

- Representation *vs.* non-representation debate
- Replacement hypothesis on embodiment
 - Cognitive system does not have to represent anything because all the information a cognitive system needs is already immediately accessible as a consequence of its non-stop real-time sensorimotor interaction with the world around it
- Arguments against:
 - examples are not “representation hungry”
 - they do not involve situations where the cognitive agent has to act on the basis of knowledge which is *not presently available to it*

Representation and Anti-Representation

Representation and Sharing Knowledge

- Cognitivist approach
 - Representation of the world is a direct **one-to-one mapping** between an (symbolic) **internal state** and its counterpart in the **real world**
 - This mapping is established by perceptual processes
 - Assumes **the things we perceive in the world *are* just as we perceive them**
 - **All other cognitive systems perceive the world in the same way**
 - **Sharing knowledge among cognitive systems poses no problems in principle**

Representation and Anti-Representation

Representation and Sharing Knowledge

- Emergent approach

- Perceptions and understanding are fundamentally linked to the manner in which you interact with that reality
- Your perceptions, **and therefore your representations of what it is you are perceiving**, are shaped by your actions and the range of possible actions you can perform
- It is **not** possible for a human designer to **implant knowledge directly** into an artificial cognitive system
- Knowledge must be acquired by an embodied cognitive agent by **learning**

Representation and Anti-Representation

What Qualifies as a Representation?

- One view: **any stable state of a cognitive system**, and of its memory in particular, **that correlates with events in the world is a representation**
- Not sufficient condition according to some experts in the field

Representation and Anti-Representation

What Qualifies as a Representation?

- To qualify as a representation these states — these “stand-ins” for the things in the world that are not immediately accessible to the cognitive agent — must also
 - be **used** for some purpose or **function**
 - be **generally available** for such use by the cognitive system
- A representation must play an active causal role in generating the system’s behaviour

Representation and Anti-Representation

Weak and Strong Representation

- **Weak** representations correspond to events that are **currently accessible** by our senses
- **Strong** representation correspond **to those that are not** (e.g. objects that are out of sight or that we saw previously)
 - Required in circumstances where the events to be represented might
 - no longer be present
 - might not even exist
 - might be **counter-factual** ... the opposite of affairs as they appear to be

Representation and Anti-Representation

Radical Constructivism

- The constructive aspect of enactivism is referred to as
 - **Constructivism** [Riegler 2005]
 - **Radical constructivism** [von Glaserfeld 1996, 1996]
- **Radical**: emphasizes that the principles of constructivism have to be applied at every level we chose to describe a cognitive system

Representation and Anti-Representation

Radical Constructivism

- (Radical) constructivism rejects representationalism
 - In the sense that representationalism assumes an external world to which cognitive agents **have direct access and can represent**
- Constructivism does allow for knowledge
 - The result of an active **process of construction** whereby the cognitive agent determines through its structural coupling with its environment **what matters** for its survival and what doesn't
 - Sense-making (enaction)
 - Model generation (computational modelling)

The Symbol Grounding Problem

- Assuming a cognitive system has some form of symbolic representation of the world around it
 - i.e. some set of tokens that denote objects in the agent's world

- **How does the representation, i.e. the symbols, acquire meaning?**

How do purely symbolic representations acquire semantic content?

- These apparently innocent questions are made **difficult** by the fact that (physical) **symbol systems are governed by purely syntactic processes**

The Symbol Grounding Problem

- Physical symbol systems
 - Atomic symbols
 - Strings of symbols
 - Symbol-based rules that define the manipulation and recombination of symbols and strings of symbols
- Defined in terms that make no reference to what these symbols mean

The Symbol Grounding Problem

- But they are all “**semantically interpretable**”
 - **the syntax can be assigned a semantic meaning** so that symbols and strings of symbols can represent objects, events, or concepts, and describe them or stand in for them
- The problem **is how to assign this meaning**
- This is the **symbol grounding problem** [Harnad 1990]

The Symbol Grounding Problem

- Symbolic representations have to be **grounded bottom-up** in **non-symbolic representations** of two kinds:
 - **Iconic** representations
 - Derived directly from sensory data (e.g. visual imagery, motor imagery)
 - Allow you to discriminate between different objects
 - **Categorical** representations
 - Based on the output of both **learned** and **innate** processes that detect **invariant features** of object & event categories from these sensory data (e.g. **object affordances**, **repeated behaviours** ... **food is nice; dogs bite**)

The Symbol Grounding Problem

- Higher-order symbolic representations can then be derived from these elementary symbols
- Both types of representation are non-symbolic
- **A non-symbolic process is required to learn the invariances and thereby form the categories**

The Symbol Grounding Problem

- Usually, we use some form of connectionist approach
 - Create the mapping
 - Form the categorical representation
- As a consequence, according to this argument, **a grounded symbol system is a hybrid system: a combination of symbolic and emergent approaches (h not H)**

The Symbol Grounding Problem

- A restricted form is the **symbol anchoring problem**
[Coradeschi and Saffiotti 2003]
- Concerned only with artificial systems

The Symbol Grounding Problem

- Establish a relationship between a symbolic label denoting some object and the sensory perception of that object
- Maintaining that relationship over extended periods of time, even when that object cannot be seen
- Only concerned only with grounding physical objects
 - Doesn't address the grounding of abstract concepts (war or peace)

The Symbol Grounding Problem

- Not everyone agrees with this view of the symbol grounding problem
- An alternative viewpoint
 - Internal symbolic representations are the result of ontogenetic development
 - They are **tethered** to the world through sensory perception rather than being *grounded*
 - **Symbol tethering** is also known as **Symbol attachment**
 - [Sloman and Chappell 2005]

The Symbol Grounding Problem

An alternative viewpoint

- Symbols don't derive directly from the sensory data (as they do in symbol grounding), they derive from development
 - The process of developing new items of knowledge that are specific to the embodiment of the agent in question

The Symbol Grounding Problem

An alternative viewpoint

- **Symbol grounding**: a **cognitivist approach**
- **Symbol tethering** is **more neutral**: makes no strong claims about
 - the relationship between world and representation
 - the necessary uniqueness of these representations
- Representations that **denote** objects *vs.* those that **connote** objects
 - absolute designation *vs.*
 - convenient association

Recommended Reading

Vernon, D. *Artificial Cognitive Systems – A Primer*, MIT Press, 2014;
Chapter 8.